

GROW

WHITE PAPER

Responsible AI in mental health: The Grow Therapy model for integrating AI into care

Authors: Kevin Ramotar, Psy.D., Matt Scult, Ph.D., Manoj Kanagaraj, MD, Bre Vergess, MPP, Alan Ni

Acknowledgements: Kelsey McAlister, Ph.D., MS, Alex Roe, Victoria Tan, J.D., Jacinta M. Jiménez, Psy.D., Ethan Goh, MD, Walter Sinnott-Armstrong, Ph.D., Valerie Hoffman, Ph.D., MPH, Steve Duke

Contents

- 01 Executive summary
- 02 The AI shift is already happening
- 03 Where AI shows promise & where it falls short
- 04 The Grow Therapy approach:
responsible AI that extends care
- 05 Grow Therapy's AI in action
- 06 Conclusion
- 07 Grow Therapy's advisory board
- 08 Reference list

01 Executive summary

For most people in therapy, care is concentrated in a single weekly hour. What happens in the other 167 hours is largely left to the client. Increasingly, they are filling that space with AI. Among U.S. adults who use large language model (LLM) chatbots, nearly half report using them for mental health needs, and more than a third say they find them more helpful than traditional therapy. Most are doing so outside of any clinical relationship, without provider knowledge or oversight.

But most of these tools were not designed for clinical use. They lack access to client history, operate without clinical guardrails, and are not connected to a broader system of care. When clients use them as a substitute for treatment, or as a between-session resource without provider involvement, there is no mechanism to ensure safety, accuracy, or alignment with their care.

Grow Therapy (Grow) takes a different approach: creating purpose-built AI tools that operate within the clinical relationship, with providers at the center. The model is guided by three core pillars: improving client outcomes, ensuring strong oversight and safety, and protecting the therapeutic relationship. It is grounded in clinical science and established AI governance frameworks. In practice, this includes tools like “coach,” Grow’s conversational AI tool that helps clients reflect, process challenges, and apply evidence-based skills between sessions while remaining connected to their provider. Early data show strong adoption and engagement, with nearly 80,000 clients using coach. Grow’s AI also supports providers through tools like “Session Insights,” which generates structured post-session summaries that providers can share directly with clients to reinforce continuity between visits, with more than one million generated as of May 2026.

Grow’s approach demonstrates how AI can be designed to operate within care, extending support beyond the session while maintaining clinical integrity, safety, and human connection.

02 The AI shift is already happening

More than one in five U.S. adults are in need of mental health support, but only about 50% receive care in a given year and over 120 million people live in mental health provider shortage areas ^[1]. That supply-demand mismatch helps explain why so many people are turning to general-purpose AI tools for support. But the fact that people find these tools accessible does not mean they are adequate. Most were not designed for clinical use, yet they are increasingly being used as substitutes for care.

This is not a hypothetical shift. People are already using AI to reflect, process emotions, and navigate challenges in real time, often without any clinical oversight. Whether these tools help or harm may depend less on the technology itself and more on how it is designed, who oversees it, and whether it is connected to a broader system of care.

23.1%

U.S. adults are in need of
mental health support

120M

People live in mental health
provider shortage areas



03 Where AI shows promise & where it falls short

AI can expand access to mental health support by offering immediate, accessible engagement and a space for reflection, particularly between therapy sessions. Research suggests that people may share more openly with conversational systems they perceive as nonjudgmental, and some studies have shown the potential for AI tools to support elements of the therapeutic alliance ^[2]. But these benefits depend on how AI is designed and used.

When AI is used for mental health outside of a clinical context, it introduces real safety risks. For example, recent reporting has documented severe adverse events tied to the use of general-purpose chatbots for mental health needs. In one widely covered case, a wrongful-death lawsuit alleges that a 14-year-old became intensely attached to a chatbot that reinforced suicidal ideation prior to his death^[3].

Safety concerns extend beyond individual cases. Mental health experts have raised warnings about chatbot-induced delusions, reinforcement of harmful thinking patterns, and failure to recognize or appropriately respond to crisis situations ^[4, 5]. Policymakers have taken action, with several states moving to restrict or ban certain types of AI chatbots in mental health care ^[6, 7], yet many of the most widely used general-purpose chatbots remain largely outside the scope of existing regulation.

These risks reflect how the tools are built. General-purpose AI systems were not designed for clinical mental health care. They lack access to client history, operate without accountability structures, and are not embedded within a broader system of care. How well AI performs in clinical settings depends in large part on whether it has access to that context and whether a provider is positioned to act on what it surfaces.

It is worth noting that the landscape is evolving. Some general-purpose tools now allow users to upload medical records, and many companies have begun building internal accountability structures. However, Grow's position is that AI is most safely and effectively applied when it operates within a care delivery network where a licensed provider can curate its use, stay informed by its interactions, and remain in the driver's seat of the clinical experience.

This does not mean AI cannot play a role in mental health care, but rather that its role must be intentionally designed and connected to existing structures that deliver quality and effective care.

Grow's Position

AI is most safely and effectively applied when it operates within a care delivery network where a licensed provider can curate its use, stay informed by its interactions, and remain in the driver's seat of the clinical experience.

04 The Grow Therapy approach: responsible AI that extends care

Grow was founded to offer therapy that’s accessible, effective, and grounded in meaningful human connection. Our AI strategy is central to achieving these goals, developed in close collaboration with clinicians, providers, and advisors rather than in isolation. The approach is grounded in both clinical science and established AI governance frameworks (see Table 1; [8, 9]). It draws on research showing the importance of therapeutic alliance and between-session engagement in driving meaningful client outcomes, and aligns with guidance from organizations such as the World Health Organization (WHO), the National Institute of Standards and Technology (NIST), and the Office of the National Coordinator for Health IT (ONC). These frameworks consistently emphasize the importance of human oversight, safety, and continuous evaluation.

PRINCIPLE	ALIGNED ORGANIZATIONS	GROW THERAPY APPROACH
Human oversight	WHO, NIST, ONC	Providers remain central to care, with visibility, control, and final decision-making authority over AI-supported interactions
Safety and risk management	WHO, NIST, Health Canada, MHRA	AI systems include guardrails, automated safety monitoring, and escalation pathways so high-need situations are routed to human providers
Clinical appropriateness	WHO, ONC	AI is designed to support care delivery, not diagnose, treat, or replace providers, and operates within clearly defined clinical boundaries
Transparency and user awareness	WHO, ONC, NIST	AI features are clearly communicated, optional, and designed with explicit boundaries so users understand their role and limitations
Data protection and privacy	WHO, NIST, HIPAA	Sensitive health data is protected through privacy-by-design practices, secure infrastructure, and strict access controls
Continuous evaluation and improvement	NIST, FDA, WHO	AI systems are continuously monitored through clinical QA, benchmarking, and iterative improvement processes to refine safety and effectiveness over time

Grow's AI principles are guided by three core pillars: 1) Client Outcomes, 2) Strong Oversight, and 3) Protecting the Therapeutic Relationship.

01

Client outcomes

AI should improve care, not just make it more efficient. AI is used to support better outcomes by helping clients stay engaged between sessions, making it easier for providers to deliver high-quality care, and expanding access where there are gaps.

AI is designed to support care and not make decisions. It operates within clear clinical boundaries, with guardrails and escalation pathways in place so that higher-risk situations are routed to licensed providers.

02

Strong oversight

AI must operate within clear systems of oversight, safety, and accountability. Grow's AI operates under a defined AI governance structure that includes training for providers, input from a clinical product team of licensed psychologists, continuous quality assurance processes, and guidance from external AI advisors.

In practice, coach illustrates what this looks like operationally: automated safety guardrails that monitor conversations in real time, an LLM-based evaluation system that assesses clinical quality at scale, and ongoing review of both randomly sampled and flagged interactions by Grow's internal clinical team. When concerns arise, treating providers are alerted and can adjust how AI tools are used for specific clients. Grow also actively tracks emerging U.S. policy and broader industry commitments, evolving its approach alongside new standards and expectations.

03

Protect the therapeutic relationship

Mental health care depends on trust, empathy, and the relationship between client and provider. AI should strengthen this relationship, not replace it.

Grow's AI is designed to support how care actually happens by helping clients reflect between sessions, making it easier to communicate what they are experiencing, and giving providers better visibility into those experiences. It is also designed to respond to tone, timing, and context, so interactions feel supportive without attempting to replicate therapy itself.

This approach helps sustain the therapeutic relationship by reducing burden on providers and reinforcing the work happening between sessions. It also recognizes that care is not one-size-fits-all, and that experiences of distress, communication, and support can vary across individuals.

Together, these principles guide how Grow evolves and uses its AI.

05 Grow Therapy's AI in action

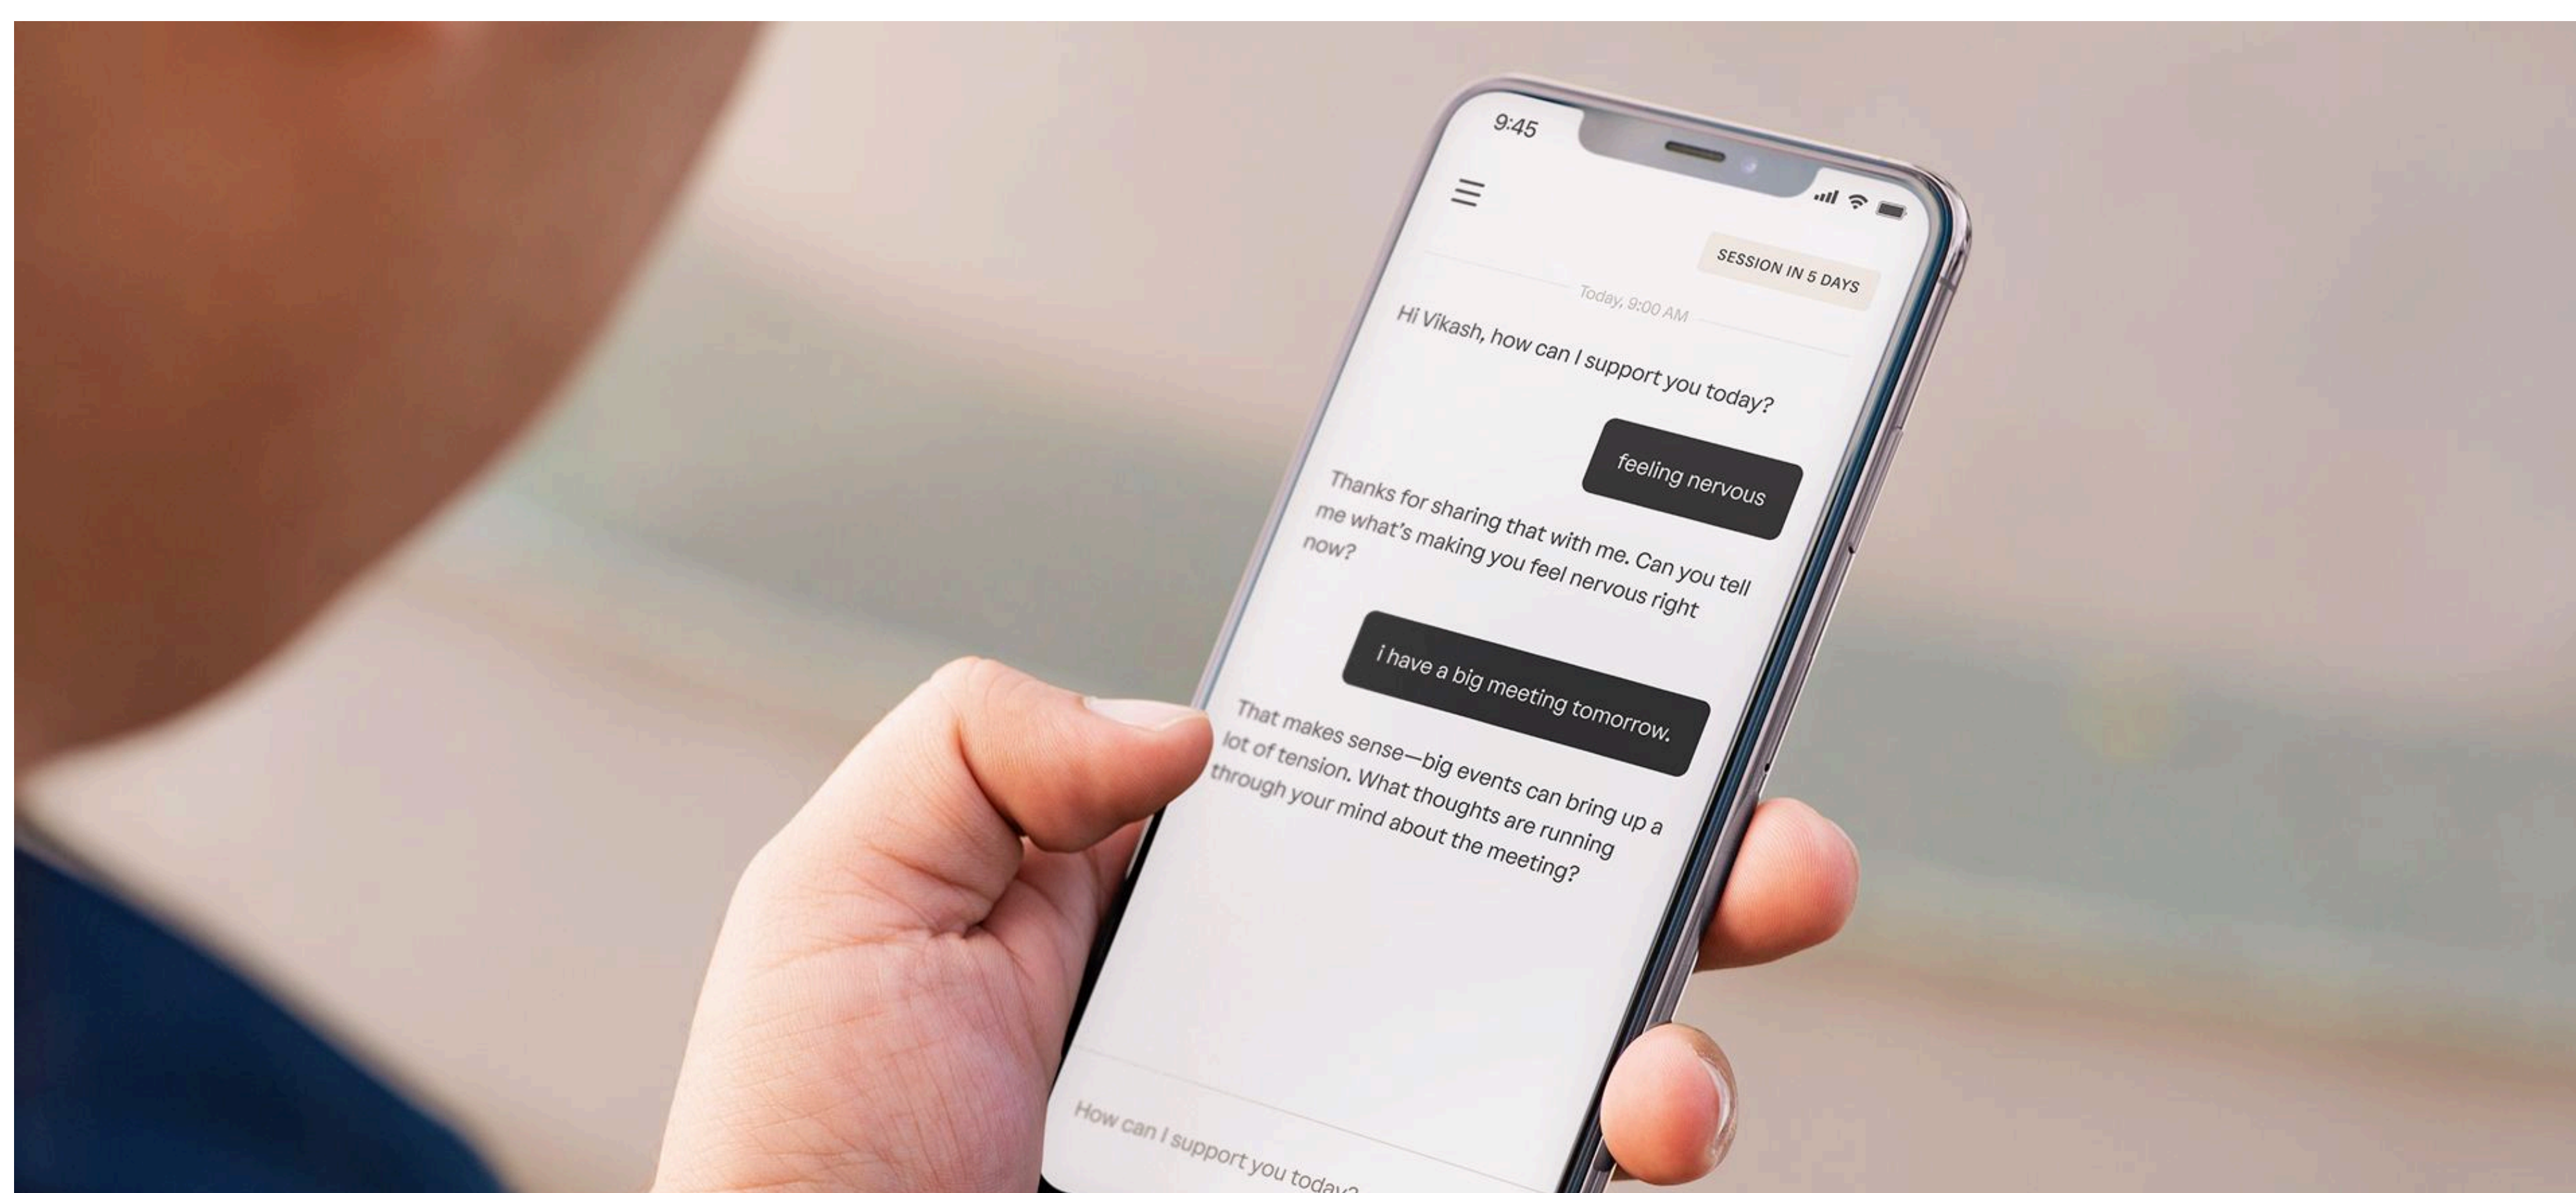
At Grow, AI is embedded into workflows that support both clients and providers across the course of care.

Extending Care Between Sessions: Coach in Practice

Most of the work in therapy happens outside the session. Clients are learning to apply new skills, sitting with difficult emotions, and navigating situations their provider will not hear about until the next appointment. Providers have historically had little visibility into this between-session time.

Coach is Grow's conversational AI tool designed to address that gap by helping clients reflect, process challenges, and practice skills between sessions. It is not a standalone product but designed to complement the clinical relationship. Clients can use coach to work through situations in real time, reinforce evidence-based skills introduced in therapy, and capture topics to bring back to their provider. Providers can review coach interactions, use them to inform session priorities, and decide whether coach is appropriate for a given client. Clinical judgment and decision-making authority remain with the treating provider throughout.

This design reflects a deliberate clinical philosophy. Coach does not diagnose, provide treatment, or position itself as a therapeutic substitute. It operates within boundaries that are explicit to users from the start and reinforced throughout. When potential safety concerns are detected, interactions are paused, crisis resources are surfaced, and providers are notified. Grow's internal clinical team reviews a portion of all interactions, alongside an automated evaluation system that assesses clinical quality and safety at scale. These oversight mechanisms are integrated into how the product operates rather than layered on after the fact.



Early data suggest the model is working. In its first five months, nearly 80,000 clients have used coach. Language patterns suggest clients feel better by the end in a majority of interactions. These are early signals, not clinical proof, but they are consistent with the hypothesis that between-session support, when integrated into a care relationship rather than offered in isolation, can reinforce rather than replace the work clients do with their therapist.

What these numbers reflect, more than adoption, is that providers and clients are using coach the way it was designed: as part of overall care.

Supporting Providers: AI in the Care Workflow

Coach supports clients in the space between sessions. Grow's provider-facing AI tools are designed to work alongside providers throughout the care workflow, not simply to reduce administrative work, but to surface clinical insight, support continuity of care, and keep providers focused on what only they can do.

Grow's AI notetaker is designed to reduce that burden at multiple points in the care workflow. During sessions, it captures key themes and details so that providers can stay focused on the clinical interaction rather than splitting attention between listening and note-taking. After sessions, it generates structured summaries and next steps, referred to as "Session Insights," which providers can review and share with clients to reinforce progress and maintain continuity between visits. Over time, these tools also help surface patterns in how clients are progressing, giving providers a more complete picture to inform clinical decision-making.

These tools are already integrated at scale across Grow's provider network. More than 1.5M+ Session Insights have been delivered to clients, and providers report reductions of up to 40% in time spent on intake notes and up to 66% for progress note documentation. These tools reflect the same principle that guides Grow's AI strategy more broadly: AI should reduce the work around care while keeping clinical judgment and decision-making authority with the provider.

Reduction in time spent on intake notes



Reduction in time spent on progress notes



More than 1.5M+ session insights have been delivered to clients

06 Conclusion

AI is already part of how people manage their mental health. The question is not whether AI belongs in this space but how it should be built, governed, and integrated into care.

Grow's approach offers a model: purpose-built AI that operates within the clinical relationship, with providers retaining decision-making authority and multiple layers of oversight working together to monitor safety and quality. Early data from coach and Grow's provider-facing AI tools suggest this model is gaining traction with both clients and providers, though Grow treats these as early signals that warrant continued evaluation rather than proof of concept.

What generalizes beyond Grow is the underlying principle. AI in mental health care should be held to high standards of oversight, accountability, and continuous improvement. As the field evolves, the organizations that get this right will be the ones that treat responsible AI not as a set of principles on paper but as an operational commitment embedded in how their products work every day.

07 Grow's AI advisory panel

Grow believes independent experts serving as accountability partners are critical to setting and maintaining a high bar for responsible AI in mental health care. To support that commitment, the company established an external Advisory Panel to help inform Grow's approach to responsible AI in mental health care more broadly. The panel brings together expertise spanning psychology, clinical AI research, mental health epidemiology, ethics, and consumer insights, providing interdisciplinary guidance grounded in both scientific rigor and compassion for the people Grow serves. Their counsel supports Grow's approach to safety, evaluation, and continuous improvement as the AI coach evolves.



Jacinta M. Jiménez, Psy.D
Clinical Psychology

Dr. Jacinta M. Jiménez is a Stanford-trained licensed clinical psychologist, board-certified coach (BCC), and leadership strategist whose work bridges behavioral science, technology, and human resilience. For more than 15 years, she has designed evidence-based frameworks that help people and organizations thrive under pressure while staying grounded in ethics and purpose.

As Vice President of Coaching Innovation at BetterUp, Dr. Jiménez played a key role in architecting and scaling the company from its earliest stages into a global leader in digital coaching, serving Fortune 500 enterprises worldwide. There, she founded and chaired the company's first Ethics Committee, establishing governance standards for the responsible use of behavioral science and technology. Earlier in her career, she contributed to the development of one of the first clinically informed mobile mental-health apps for the National Center for PTSD. An award-winning author of *The Burnout Fix* (McGraw-Hill), her work has been featured in *Harvard Business Review*, *Fast Company*, and *Business Insider*. She has spoken at NASA, Google, and Columbia University's Zuckerman Institute, equipping leaders with tools for ethical and resilient leadership.

Dr. Jiménez also serves on the Executive Board of Abroad.io and advises organizations on building human-centered systems that advance discernment, adaptability, and connection in the age of AI.



Ethan Goh, MD
Clinical AI Research and Scientific Evaluation

Dr. Ethan Goh is Executive Director at ARISE, a Stanford-Harvard research network focused on how frontier AI models are safely translated into real world care. His research is supported by the Macy Foundation and the Gordon & Betty Moore Foundation, and has been covered in *The New York Times*, *The Washington Post*, and *CNN*.

Dr. Goh is a Founding Editorial Board member and Associate Editor at *BMJ Digital Health & AI*, and an external reviewer for multi-year research grants.



Walter Sinnott-Armstrong is Chauncey Stillman Professor of Practical Ethics in the Department of Philosophy and the Kenan Institute for Ethics at Duke University. He has secondary appointments in the Law School and the Department of Psychology and Neuroscience, and he is core faculty in the Duke Center for Cognitive Neuroscience, the Duke Institute for Brain Sciences, and the Duke Center for Interdisciplinary Decision Sciences.

He was born in Memphis, Tennessee, and attended the Hotchkiss School, Amherst College (B.A. 1977), and Yale University (Ph.D. 1982). He taught at Dartmouth College from 1981 to 2009 and moved to Duke University in 2010.

Walter Sinnott-Armstrong, PhD
Ethics Scholar



Valerie is passionate about improving the state of solutions developed to help those affected by mental health challenges. As the current CEO and Founder of Pogo Research, she works cross-functionally with digital behavioral health companies to understand their value by suggesting valid data assessments and variables to collect and analyze, as well as studies to conduct that will provide insights into how their solutions improve access, efficacy, satisfaction, equity, and cost.

She has authored over 100 publications in peer-reviewed journals and presented at over 100 national scientific conferences. She was the former VP of Woebot Health and Chief Research Officer at Meru Health and a psychiatric epidemiologist by training, earning an MPH from Yale University and a PhD from Johns Hopkins Bloomberg School of Public Health, where she was an NIMH predoctoral fellow.

Valerie Hoffman, PhD, MPH
Mental Health Epidemiology



Steve Duke is the founder of Hemingway, an organization supporting behavioral health innovators. He is also the author of The Hemingway Report, a newsletter providing industry analysis to thousands of behavioral health leaders. His work focuses on using data to understand the intersection of technology, markets, and clinical innovation, and translating that understanding into actionable insights. Before Hemingway, Steve worked as a management consultant at McKinsey and led growth teams at two tech unicorns (LetsGetChecked & Wayflyer).

Steve Duke
Insights and Strategy

08 Reference list

- [1] Rousmaniere T, Zhang Y, Li X, et al. Large language models as mental health resources: Patterns of use in the United States. *Pract Innov*. Epub ahead of print 2025. DOI: 10.1037/pri0000292.
- [2] Fitzpatrick KK, Darcy A, Vierhile M. Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Ment Health* 2017; 4: e7785.
- [3] Silicon Valley AI company sued over 14-year-old's suicide. ABC7 San Francisco, <https://abc7news.com/post/silicon-valley-based-company-character-ai-sued-florida-14-year-olds-suicide/15461471/> (2024, accessed 8 April 2026).
- [4] Preliminary Report on Dangers of AI Chatbots | Psychiatric Times, https://www.psychiatrytimes.com/view/preliminary-report-on-dangers-of-ai-chatbots?utm_source=chatgpt.com (accessed 3 April 2026).
- [5] "AI Psychosis": How ChatGPT Amplifies Delusions & Triggers Psychosis. *Psychiatry & Psychotherapy Podcast*, <https://www.psychiatrypodcast.com/psychiatry-psychotherapy-podcast/episode-253-ai-psychosis-emerging-cases-of-delusion-amplification-associated-with-chatgpt-and-llm-chatbot> (accessed 3 April 2026).
- [6] 10/03/2025: States struggle to regulate AI chatbots for mental health therapy amid rising need for care | News, <https://www.healthpolicyohio.org/health-policy-news/2025/10/03/states-struggle-to-regulate-ai-chatbots-for-mental-health-therapy-amid-rising-need-for-care> (accessed 3 April 2026).
- [7] © States set new rules for AI chatbots in mental health care | LinkedIn, <https://www.linkedin.com/pulse/states-set-new-rules-ai-chatbots-mental-health-care-mathematica--wpkie/> (accessed 3 April 2026).
- [8] Mausbach BT, Moore R, Roesch S, et al. The Relationship Between Homework Compliance and Therapy Outcomes: An Updated Meta-Analysis. *Cogn Ther Res* 2010; 34: 429–438.
- [9] Autio C, Schwartz R, Dunietz J, et al. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST, <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence> (2024, accessed 3 April 2026).